IDEAL TRIANGULATIONS OF PSEUDO-ANOSOV MAPPING TORI

IAN AGOL

Dedicated to Bus Jaco on the occasion of his 70th birthday

ABSTRACT. We show how to construct an ideal triangulation of a mapping torus of a pseudo-Anosov map punctured along the singular fibers. This gives rise to a new conjugacy invariant of mapping classes, and a new proof of a theorem of Farb-Leininger-Margalit. The approach in this paper is based on ideas of Hamenstadt.

1. INTRODUCTION

Recently, Farb, Leininger and Margalit [2] proved that the mapping tori of small dilatation mapping classes of closed surfaces are obtained by Dehn filling on finitely many cusped hyperbolic 3-manifolds. Recent results of Hamenstadt [6] suggested another approach to this theorem using splitting sequences of train tracks. In discussions with Hamenstadt, she suggested that her complex of train track splitting sequences should give a new solution to the conjugacy problem for pseudo-Anosov mapping classes, somewhat similar to work of Mosher on "circular expansion complexes" [12, 13, 14]. In this paper we give an exposition of these results from the perspective of measured train track splitting sequences. Instead of a complex of splittings, we obtain a taut ideal triangulation associated to a mapping class. The fact that the reverse of splitting sequences (folding sequences) give rise to Perron-Frobenius maps of the weight spaces of train tracks enables us to give a new proof of the main result of [2].

What Hamenstadt proved is that if one is given a minimal lamination $\mathcal{L}$ on a surface, there is a complex of splitting sequences of train tracks carrrying $\mathcal{L}$ which forms a CAT(0) cube complex in a natural way, with vertices corresponding to train tracks carrying the lamination $\mathcal{L}$, directed edges corresponding to $\mathcal{L}$-splits of the train tracks, and cubes corresponding to commuting collections of $\mathcal{L}$-splits [6]. This is analogous to the complex constructed by Mosher [14], however we remark that the definition of "splitting sequence" in that manuscript is distinct from the usage in this paper (Mosher allows shifts/slides as well). If one takes the stable lamination associated to a pseudo-Anosov map φ, then one may find a bi-infinite cube complex of train tracks which is invariant under φ, and this gives rise to an invariant of the conjugacy class of φ. Instead of considering all sequences of splittings in this paper, we consider "maximal splitting" sequences, where the train tracks have a measure, and the splittings occur at the branches of the train track of maximal weight. We prove that these sequences are eventually periodic for a pseudo-Ansov stable lamination in Section 3 (this is somewhat analogous to the approach of [15, 17] and [14,

Lemma 10.2.6]). These sequences give something like a continued fraction expansion for pA mapping classes, analogous to the case of Anosov maps of the 2-torus (Mosher observed the analogy between continued fractions and his circular expansion complexes [14]). This approach allows us to reprove the main result of [2] in Section 6. We also obtain a layered ideal triangulation of the mapping torus punctured at the singular fibers, from Whitehead moves on the triangulations dual to the train tracks in the splitting sequence in Section 4. The natural structure of these triangulations is a taut ideal triangulation, introduced by Lackenby [10].

2. Definitions

We review some background definitions and establish some notation in this section.

Let $\Sigma = \Sigma_{g,n}$ be an orientable surface of genus g with n punctures, and let $\varphi : \Sigma \to \Sigma$ be a homeomorphism. By the Nielsen-Thurston classification of mapping classes, φ is conjugate to either an element of finite order, a reducible mapping class, or to a pseudo-Anosov map [23, 20]. In the third case, there are the stable and unstable measured laminations $\mathcal{L}^s$ and $\mathcal{L}^u \in \mathcal{ML}(\Sigma)$ associated to φ such that for any curve $c \subset \Sigma$, $[\varphi^n(c)] \to [\mathcal{L}^s] \in \mathcal{PL}(\Sigma)$, and $[\varphi^{-n}(c)] \to [\mathcal{L}^u] \in \mathcal{PL}(\Sigma)$ (see [22, Chapter 8] for the notation). Moreover, $\mathcal{L}^s$ and $\mathcal{L}^u$ meet every essential closed curve in Σ which is not isotopic into a neighborhood of a puncture. Further, there exists a dilatation $\lambda(\varphi) \in (1, \infty)$, $\varphi(\mathcal{L}^s) = \lambda(\varphi)\mathcal{L}^s$ and $\varphi(\mathcal{L}^u) = \lambda(\varphi)^{-1}\mathcal{L}^u$, up to isotopy. These laminations are only unique up to scaling and isotopy. If we choose a complete finite area hyperbolic metric on Σ, then we may isotope the measured laminations $\mathcal{L}^s$ and $\mathcal{L}^u$ to have totally geodesic leaves. Each complementary region will be isometric to an ideal polygon with at least three sides, or to a punctured disk with at least one puncture on the boundary. If we puncture Σ at a point in each ideal polygon complementary region, we obtain a surface Σ°, and we have a well-defined (up to isotopy) restriction map $\varphi^\circ = \varphi_{|\Sigma^\circ} : \Sigma^\circ \to \Sigma^\circ$. Moreover, $\lambda(\varphi^\circ) = \lambda(\varphi)$, and we may assume $[\varphi^\circ(\mathcal{L}^{u,s})] = [\mathcal{L}^{u,s}] \in \mathcal{PL}(\Sigma^\circ)$.

To encode a lamination combinatorially, we use train tracks (after Thurston [22, Ch. 8]). A train track $\tau \subset \Sigma$ is a 1-complex with trivalent vertices which is locally modeled on a switch (see Figure 1(a)) and satisfying some extra conditions. Edges of the train track are called *branches*, and vertices are called *switches*. Each branch $e \subset \tau$ is C^1 embedded, and at each switch of τ, there is a well-defined tangent space to the branches coming into the switch. A *half-branch* is an end of the interior of a branch. Each switch is in the closure of three half-branches. A *large half-branch* is the end of a half-branch meeting a switch on the the side of the tangent space to the switch opposite of the other two half-branches incident with the switch, and a *small half-branch* is a half-branch meeting the switch on the same side of the tangent space to the switch as another branch (in Figure 1(b), the half-branch labelled a is large, whereas the half-branches labeled b and c are small). For each component R of $\Sigma - \tau$, the boundary ∂R is a piecewise smooth curve. A non-smooth point of ∂R is a *cusp*. For each component R of $\Sigma - \tau$, the double of R along ∂R with the

cusps removed must have negative Euler characteristic. We will follow the notation and conventions of Penner-Harer [18]. A *large branch* is a branch with both ends being large half-branches, and a *small branch* has both ends small half-branches. Otherwise, a branch is *mixed*. When all of the complementary regions of the train track τ are punctured disks, then it is convenient to label the branches of the track by the isotopy class of the arc type of an ideal edge dual to the branch. This gives an unambiguous way of labeling branches in differing train tracks on the same surface.
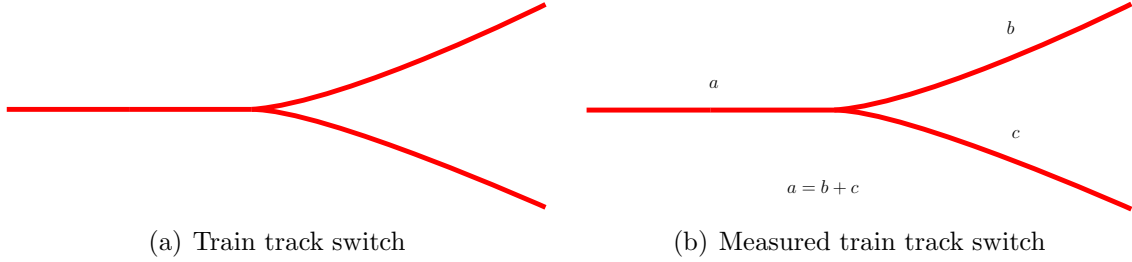

(a) Train track switch   (b) Measured train track switch

FIGURE 1.   Train track switches

A *measured train track* is a train track τ together with a transverse measure μ which is a function assigning a weight to each edge of τ. We will denote a measured train track as a pair (τ, μ). At each switch of τ with incoming edges b, c and outgoing edge a, the weights of the adjacent edges must satisfy $\mu(a) = \mu(b) + \mu(c)$ (Figure 1(b)). In some pictures, we will abuse notation and say $a = b + c$ when we really mean $\mu(a) = \mu(b) + \mu(c)$. This is especially convenient when the train track changes, but most of the labelled edges do not. The *weight space* $W(\tau)$ is the convex space of all positive measures on τ. The lamination $\mathcal{L}$ is *suited to* the train track τ if there is a differentiable map $f : \Sigma \to \Sigma$ homotopic to the identity such that $f(\mathcal{L}) = \tau$, and f is non-singular on the tangent spaces to the leaves of $\mathcal{L}$. If $\mathcal{L}$ is a measured lamination and $\mathcal{L}$ is suited to τ, then $\mathcal{L}$ induces a measure μ on τ given by the measure on τ which pulls back to the transverse measure on $\mathcal{L}$ via the map f.

A *splitting* of a measured train track (τ, μ) is a move at a large branch which splits it according to the weights of the neighboring edges (see Figure 2). We use the notation $(\tau, \mu) \rightharpoonup_e (\sigma, \nu)$.
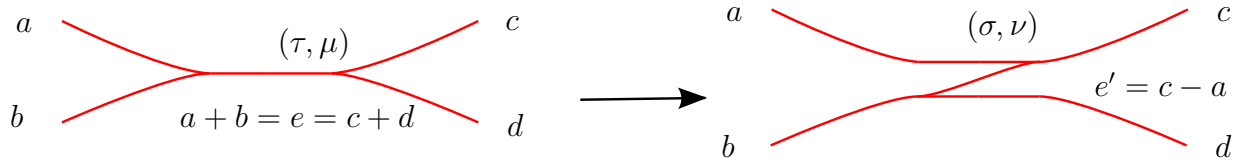

FIGURE 2.   Splitting a large branch when $a < c$.

If the measured lamination $\mathcal{L}$ is suited to (τ, μ), then $\mathcal{L}$ will also be suited to (σ, ν) when $(\tau, \mu) \rightharpoonup_e (\sigma, \nu)$. If the weights $\mu(a) = \mu(c)$, then the branch e' will not appear in the train track σ, but we do not consider this to be a split in this paper. Conversely, any measured train track (τ, μ) gives rise to a unique measured lamination $\mathcal{L}$ suited to (τ, μ).

Two other moves on train tracks do not depend on the measure. A *shift* changes the track at a mixed branch (Figure 3) (these are also called *slides* [14]). A *fold* changes the track at a small branch (Figure 4). The shift is self-inverse, whereas the fold is inverse to a split. However, the split depends on the measure, whereas the fold does not, which is why it is important to make a distinction.
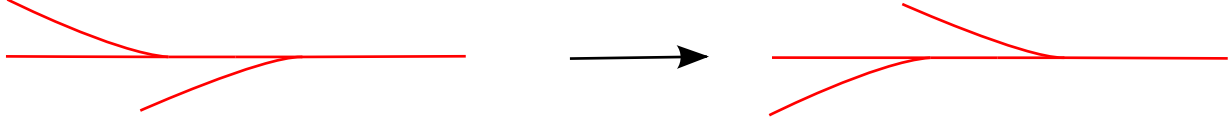

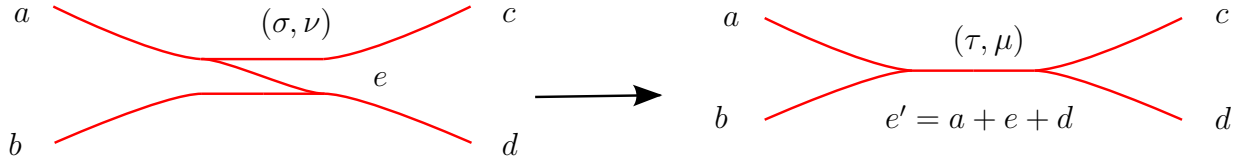FIGURE 3.   Shifting a mixed branch.

FIGURE 4.   Folding a small branch.

Two measured train tracks $(\tau, \mu), (\tau', \mu')$ are *equivalent* if there is a sequence of train tracks obtained by splits, shifts, folds and isotopies which interpolates between the two.

A *maximal splitting* of (τ, μ) is a splitting along all of the large branches of τ with maximal μ weight. We denote $(\tau_0, \mu_0) \rightharpoonup (\tau_1, \mu_1)$ if (τ_1, μ_1) is obtained from (τ_0, μ_0) by a single maximal splitting. If we have n maximal splittings $(\tau_0, \mu_0) \rightharpoonup (\tau_1, \mu_1) \cdots \rightharpoonup (\tau_n, \mu_n)$, then we denote $(\tau_0, \mu_0) \rightharpoonup^n (\tau_n, \mu_n)$, or if we don't want to specify n, we will use $\rightharpoonup^*$.

A *trainpath* on a train track τ is a C^1-immersion $\rho : [m, n] \to \tau \subset \Sigma$ which maps each interval $[k, k+1]$ ($m \le k \le n-1$) onto a branch of τ. The integer $n - m$ is then called the *length* of ρ. Each complementary region of τ is bounded by a finite number of trainpaths which either are simple closed curves or terminate at the cusps of the region. A lamination is *minimal* if every half-leaf is dense.

Lemma 2.1. *Let $\mathcal{L}$ be a minimal lamination, and let $\mathcal{L}$ be suited to (τ, μ). Let $e \subset \tau$ be a branch of τ, and let $(\tau, \mu) \rightharpoonup (\tau_1, \mu_1) \rightharpoonup (\tau_2, \mu_2) \rightharpoonup \cdots \rightharpoonup (\tau_n, \mu_n) \rightharpoonup \cdots$ be an infinite sequence of maximal splittings of (τ, μ). Then there exists n such that $(\tau_n, \mu_n) \rightharpoonup (\tau_{n+1}, \mu_{n+1})$ splits the branch e (so that $\mu_n(e)$ is maximal weight for τ_n).*

Proof. We may assume that we have a map $f : \Sigma \to \Sigma$ such that $f(\mathcal{L}) = \tau$ since $\mathcal{L}$ is suited to τ. For any half-leaf $l \subset \mathcal{L}$, since $\mathcal{L}$ is minimal, $f(l)$ must eventually cross the branch e. For each cusp of τ, there are two half-leaves of $\mathcal{L}$ which are parallel for all time (since $\mathcal{L}$ is suited to τ) whose start is adjacent to the cusp (these leaves are the ideal arcs corresponding to the cusp in the complementary ideal polygon region of $\mathcal{L}$). For each cusp c of τ, let $\rho_c : [n, m] \to \tau$ be a trainpath which is parallel to the path $f(l)$ emanating from the cusp c, such that $\rho_c([m-1, m]) = e$. Each time we split $\tau_k \rightharpoonup \tau_{k+1}$ at a branch adjacent to a cusp c, the trainpath ρ_c shrinks to $\rho_c : [n-1, m] \to \tau$. Thus, at each stage of splitting,

the total length of such cusp paths decreases by 2 (since each splitting branch is adjacent to two cusps), and thus we see that eventually we must split the branch e. $\square$

3. SPLITTING SEQUENCES

We state some basic results about measured train tracks and measured laminations.

Theorem 3.1. [22, Prop. 8.9.2], [16, Theorem 4.1] *If $\mathcal{L}$ is a measured lamination, then $\mathcal{L}$ is suited to a measured train track (τ, μ).*

Theorem 3.2. [18, Theorem 2.8.5] *Measured train tracks (τ_i, μ_i) give rise to the same measured lamination $\mathcal{L}$ (up to isotopy) if and only if they are equivalent.*

Together, these theorems imply that equivalence classes of measured train tracks are in one-to-one correspondence to measured laminations. Equivalence classes of measured train tracks are generated by splitting, shifting, and folding. The next theorem implies that shifting is not needed.

Theorem 3.3. [18, Theorem 2.3.1] *If (τ_i, μ_i) are equivalent positively measured train tracks, then there is a train track (τ, μ) such that (τ_i, μ_i) splits (and isotopes) to (τ, μ) for $i = 1, 2$.*

This theorem is proven by demonstrating that two train tracks which are equivalent by a split, shift, or fold, have common splittings (having a common splitting is not hard to show is an equivalence relation using the technique of Corollary 3.4). For splits and folds this is clear. For two train tracks related by a shift, one performs the same sequence of splits of maximal edges to a sequence of two train tracks related by shifts to see that once the middle incoming edge b becomes maximal, the two resulting train tracks are the same (see Figure 5).
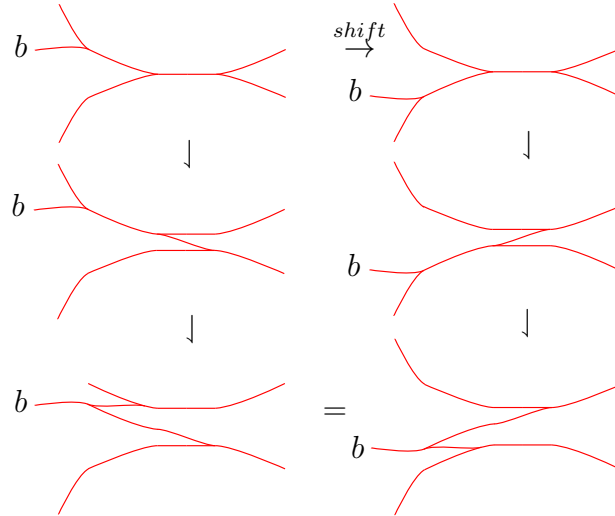

FIGURE 5. The effect of shifting then splitting.

Corollary 3.4. *Suppose $\mathcal{L}$ is suited to (τ, μ), (τ', μ'). Then there exists (τ'', μ'') such that $(\tau, \mu) \rightharpoonup^* (\tau'', \mu'')$ and $(\tau', \mu') \rightharpoonup^* (\tau'', \mu'')$.*

Proof. The property of having a common maximal split is transitive, so by Theorem 3.3, we need only assume that $(\tau, \mu) \rightharpoonup_e (\tau', \mu')$ for some large branch $e \subset \tau$. If e is the sole branch of maximal weight of (τ, μ), then we have $(\tau, \mu) \rightharpoonup (\tau', \mu') = (\tau'', \mu'')$ (a maximal split). If there are other branches of (τ, μ) of maximal weight $= \mu(e)$, then we have $(\tau, \mu) \rightharpoonup (\tau'', \mu'')$ and $(\tau', \mu') \rightharpoonup (\tau'', \mu'')$, where this maximal splitting is along the edges of τ' which have the same μ' weight as $\mu(e)$.

Otherwise, assume that $\mu(e)$ is not a maximal weight of μ. Let $(\tau, \mu) = (\tau_1, \mu_1) \rightharpoonup (\tau_2, \mu_2) \rightharpoonup \cdots$ be the maximal splitting sequence. By Lemma 2.1, there exists n such that $\mu_n(e)$ is a maximal weight of μ_n.

We claim that $(\tau_1, \mu_1) \rightharpoonup^n (\tau_{n+1}, \mu_{n+1})$ and $(\tau', \mu') \rightharpoonup^{n-1} (\tau_{n+1}, \mu_{n+1})$. We'll assume that $(\tau_i', \mu_i') \rightharpoonup (\tau_{i+1}', \mu_{i+1}')$, for $i = 2, \ldots, n$, where $\tau_2' = \tau'$. By induction, we prove that $(\tau_i, \mu_i) \rightharpoonup_e (\tau_{i+1}', \mu_{i+1}')$ for $i < n$. Once e becomes a maximal weight edge of (τ_n, μ_n), there may be other maximal edges of (τ_n, μ_n) of the same weight as $\mu_n(e)$, but we will have $(\tau_n, \mu_n) \rightharpoonup (\tau_{n+1}, \mu_{n+1}) = (\tau_{n+1}', \mu_{n+1}')$.

$$
\begin{array}{ccccccc}
(\tau_1, \mu_1) & \rightharpoonup & (\tau_2, \mu_2) & \rightharpoonup \cdots \rightharpoonup & (\tau_n, \mu_n) \\
\downarrow e & & \downarrow e & & \downarrow e \\
(\tau_2', \mu_2') & \rightharpoonup & (\tau_3', \mu_3') & \rightharpoonup \cdots \rightharpoonup & (\tau_{n+1}, \mu_{n+1})
\end{array}
$$

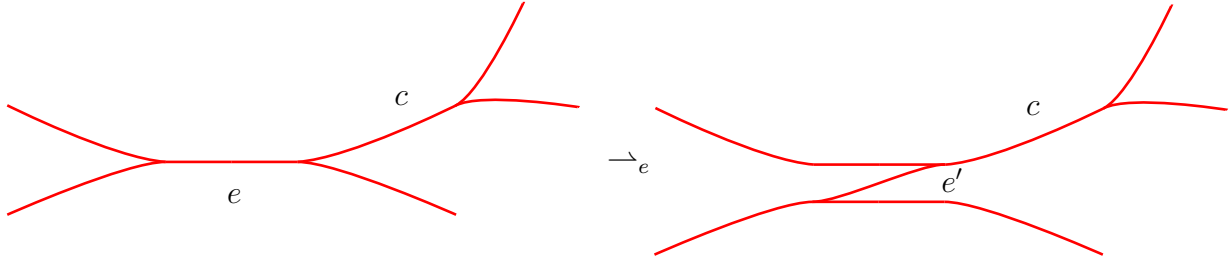Let e' be the new edge of τ_2'. To see the induction, we prove that if $\mu_i(e)$ is not maximal

FIGURE 6. Splitting a large branch to get a new large branch $c < e$.

weight, then τ_i and τ_{i+1}' agree outside of a neighborhood of e and e' consisting of the edges incident with e. The point is that after splitting e, the only new large branches c that may appear must be incident with e' (see Figure 6). However, the weight of these edges is $\mu_i(c) < \mu_i(e)$, and therefore they will not be maximal weight in $(\tau_{i+1}', \mu_{i+1}')$ if e is not maximal weight in (τ_i, μ_i). Thus, the maximal splittings of τ_i, τ_{i+1}' will occur at the same edges, which will be disjoint from the edges adjacent to e and e', respectively. So we will have $\tau_{i+1} \rightharpoonup_e \tau_{i+2}'$. $\square$

Now, we observe that if $\mathcal{L}$ is suited to (τ, μ), then it will also be suited to (τ'', μ''). The argument in Corollary 3.4 can be used to show that any sequence of splits of a measured train track may be arranged by a sequence of commuting splits to be in maximal order.

Theorem 3.5. *If $\varphi : \Sigma \to \Sigma$ is a pA map, with stable lamination $\mathcal{L}^s$, and (τ, μ) is suited to $\mathcal{L}^s$, then there exists n, m such that*

$$(\tau, \mu) \rightharpoonup^n (\tau_n, \mu_n) \rightharpoonup^m (\tau_{n+m}, \mu_{n+m}),$$

and $\tau_{n+m} = \varphi(\tau_n)$ and $\mu_{n+m} = \lambda(\varphi)^{-1}\varphi(\mu_n)$.

Proof. We have $\varphi(\mathcal{L}^s) = \lambda(\varphi)\mathcal{L}^s$. Therefore, (τ, μ) and $(\varphi(\tau), \lambda(\varphi)^{-1}\varphi(\mu))$ are equivalent measured train tracks by Theorem 3.2.

Let
$$(\tau, \mu) \rightharpoonup (\tau_1, \mu_1) \rightharpoonup (\tau_2, \mu_2) \rightharpoonup \cdots \rightharpoonup (\tau_n, \mu_n) \rightharpoonup \cdots$$
be the sequence of maximal splits. Then we have
$$(\varphi(\tau), \varphi(\mu)) \rightharpoonup (\varphi(\tau_1), \varphi(\mu_1)) \rightharpoonup \cdots \rightharpoonup (\varphi(\tau_n), \varphi(\mu_n)) \rightharpoonup \cdots$$
is also a sequence of maximal splits. Since $(\tau, \mu) \sim (\varphi(\tau), \lambda(\varphi)^{-1}\varphi(\mu))$, by Corollary 3.4 there exists n, m such that $\varphi(\tau_n) = \tau_{n+m}$ and $\varphi(\mu_n) = \lambda(\varphi)\mu_{n+m}$. □

To summarize, the splitting sequence (τ_n, μ_n) is eventually periodic, modulo the action of φ and rescaling the measure. One can recognize this periodicity combinatorially. See the example in Figure 17.

Remark 3.6. We make a remark about the kind of train tracks which can arise in such periodic splitting sequences. There can be no edge of the train track τ which is a small branch with two branches coming off of the same side, or a monogon (see Figure 7). If τ has such an edge, then it is stable under any sequence of folds, since the vertices involved in a fold can never involve the interior of the small branch of one of these excluded configurations, so that the branch always remains small. So this edge will never disappear, contradicting Lemma 2.1. Note that the isolated monogon is the same as a one-sided small branch when pulled back to the universal cover.

(a) One sided small branch   (b) Isolated monogon

FIGURE 7.   Excluded configurations

4. TRIANGULATIONS

An ideal triangulation T of a punctured surface Σ is a decomposition along ideal arcs into ideal triangles. A *Whitehead move* takes any arc in T which is adjacent to two distinct triangles, removes it, and replaces it with the other diagonal of the quadrilateral (Figure 8).

Any two ideal triangulations of Σ are related by Whitehead moves (see [10, Lemma 6] for a proof).

The *mapping torus* of $\varphi : \Sigma \to \Sigma$ is the manifold $\mathcal{T}(\varphi) = \Sigma \times [0, 1]/\{(x, 0) \sim (\varphi(x), 1)\}$. Given a pseudo-Anosov homeomorphism $\varphi : \Sigma \to \Sigma$, let $\Sigma^{\circ}_{\varphi} = \Sigma^{\circ} \subset \Sigma$ be the surface
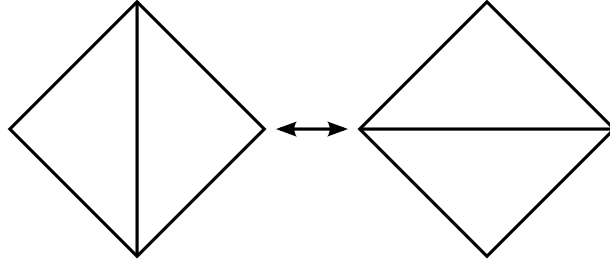
FIGURE 8. A Whitehead move

obtained by removing the singular points of the stable and unstable foliations for φ and let $\varphi^\circ : \Sigma^\circ \to \Sigma^\circ$ denote the restriction.

An *ideal triangulation* of a 3-manifold N with boundary is obtained by taking a CW complex with affinely identified tetrahedra, such that removing the vertices gives a manifold homemorphic to the interior of N. A *taut tetrahedron* is a tetrahedron such that each face is assigned a coorientation, such that two faces are pointing inwards and two are pointing outwards. Each dihedral edge of a taut tetahedron may be assigned an angle of either 0 or π, such that the sum of the angles around each corner of a vertex is π, and so that each face is co-oriented in such a way that the orientations of adjacent faces change only along an edge of angle 0. There is only one taut tetrahedron up to combinatorial equivalence (see Figure 9(a)). A *taut ideal triangulation* is an ideal triangulation of N such that each triangle is assigned a coorientation in such a way that every tetrahedron is taut, and the sum of the angles around every edge is 2π (Figure 9(b)). One may "pinch" the triangles together along an edge to obtain a branched surface, so that there are cusps along the angle zero dihedral corners, and a smooth surface between the angle π faces of the tetrahedra.
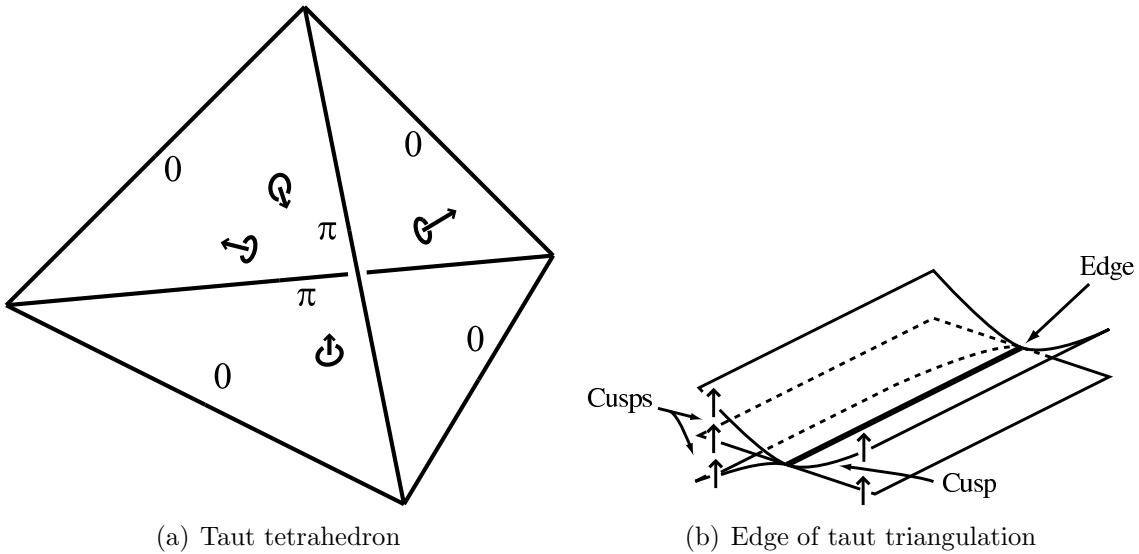

(a) Taut tetrahedron

(b) Edge of taut triangulation

FIGURE 9. Conditions for a taut ideal triangulation

If we are given $\mathcal{T}(\varphi^\circ)$ a mapping torus with ideal triangulation T of Σ°, and a periodic sequence of Whitehead moves $T \to T_1 \to \cdots \to T_m = \varphi^\circ(T)$, we may form a taut ideal triangulation of $\mathcal{T}(\varphi^\circ)$. Start with the triangulation T of Σ°. We attach a tetrahedron to Σ° along the two triangles of T which are removed under the Whitehead move to make T_1 (see Figure 10).
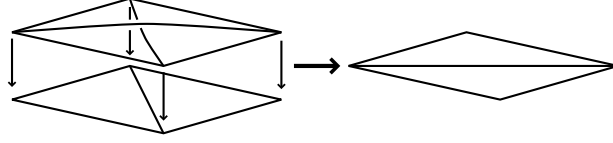

FIGURE 10. Attaching a taut tetrahedron.

The triangles are cooriented in a consistent fashion with Σ°. We repeat this process until we get T_m, which then may be glued to T by φ. The fact that φ° is pseudo-Anosov guarantees that we get a triangulation of $\mathcal{T}(\varphi^\circ)$. This triangulation is called a *layered triangulation*.

Main Construction:

We obtain a layered triangulation from the periodic sequence of train tracks given by Theorem 3.5. Each train track τ_i gives a spine of Σ°, which is dual to a unique ideal triangulation T_i of Σ°. For each split $\tau_i \rightharpoonup_e \tau_{i+1}$, one has a corresponding dual Whitehead move $T_i \to T_{i+1}$ (see Figure 11). Later on, we will also consider the reverse Whitehead move associated to a fold. It's also clear that if we have a maximal split $\tau_i \rightharpoonup \tau_{i+1}$, then
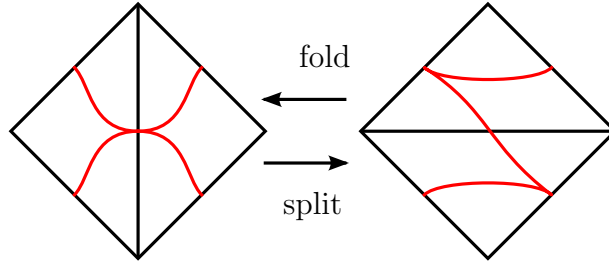

FIGURE 11. Splitting and folding

this breaks up into a sequence of splits along the maximal weight branches of τ_i, and that the order in which we attach the tetrahedra does not matter since they are attached along disjoint pairs of ideal triangles in T_i. Thus, we obtain a canonical taut ideal triangulation of $\mathcal{T}(\varphi^\circ)$ associated to the mapping class φ.

We would now like to obtain an intrinsic characterization of the taut ideal triangulations that arise from this process. We would like to be able to draw pictures of a taut ideal triangulation. In particular, we may take the preimage of the ideal triangulation in the universal cover $\mathcal{T}(\tilde{\varphi}^\circ)$, and flatten it out into the universal cover of $\tilde{\Sigma}^\circ \cong \mathbb{H}^2$, in such a way that each ideal triangle is projected so that its coorientation agrees with that of $\tilde{\Sigma}^\circ$. We will draw local pictures of triangulations circumscribed by a circle corresponding to $\partial_\infty \mathbb{H}^2$, with the convention that the coorientation is pointing toward the reader.

Definition 4.1. An edge e of a taut ideal triangulation is *left-veering* if the sequence of oriented triangles move to the left on both sides of the edge e, when viewed from the edge e and ordered by the coorientation of the triangles. Moreover, on each side of e there is at least one left-veering move. Similarly, e is right-veering if the triangles move to the right on both sides of the edge e. In particular, the degree of e must be at least 4 (see Figure 12). A taut triangulation is called *veering* if every edge is left- or right-veering.
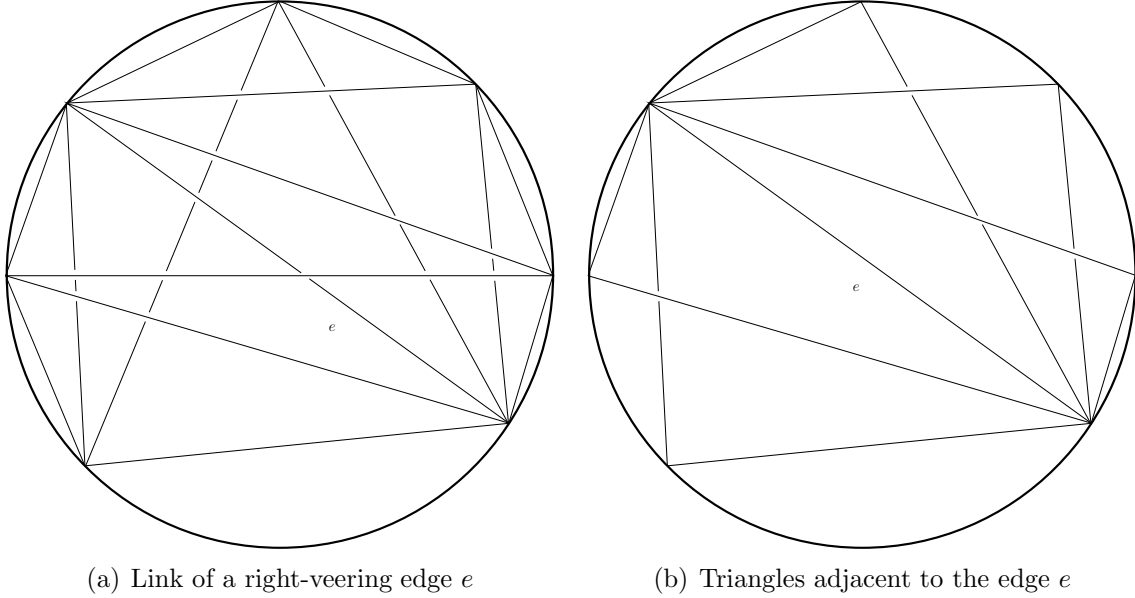

(a) Link of a right-veering edge e (b) Triangles adjacent to the edge e

FIGURE 12.

Remark: This criterion is determined purely in terms of the link of an edge in the taut ideal triangulation, so it makes sense even if the taut triangulation is not layered.

Proposition 4.2. *A taut ideal triangulation of a fibered manifold coming from Whitehead moves is associated to a periodic splitting sequence if and only if it is veering.*

Proof. The veering condition follows directly from the combinatorics of a splitting sequence, together with the condition on excluded edges in Remark 3.6. In order to see these conditions, it turns out it is necessary to work with folding sequences instead of splitting sequences. The first time an edge e appears in the sequence of Whitehead moves associated to a folding sequence, it will be dual to a large branch. In order for another Whitehead move to be applied to e, the dual branch must become small, and therefore foldings must occur involving both switches of the branch, corresponding to Whitehead moves on both sides of the edge e (viewed in the universal cover $\tilde{\Sigma}^{\circ}$). If one of these moves is left-veering and the other right-veering, then one sees that the branch dual to e becomes a one-sided small branch (Figure 13), which is excluded from a folding sequence (Figure 7). Thus, the initial Whitehead moves adjacent to e must be veering in the same direction.

The subsequent Whitehead moves on one side of e must veer in the same direction as the initial move. Initially, when the branch dual to e is large, there are two edges f, g in a triangle containing e. When we perform a Whitehead move, say along the edge g, then
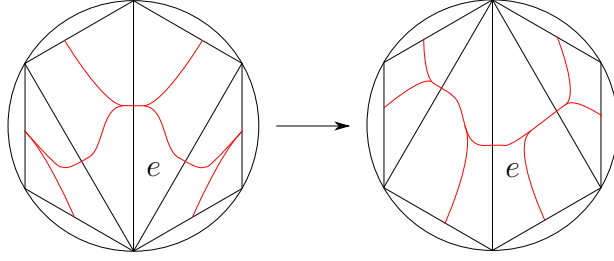
FIGURE 13. Left/right Whitehead moves give an excluded configuration

the edge g' replacing g is dual to a large branch. Thus, any subsequent move adjacent to e must be a Whitehead move dual to a fold along an edge h which is not g', and the edge h' replacing h must be dual to a large branch after such a move (see Figure 14). By induction, we see that the sequence of moves will always veer in the same direction.
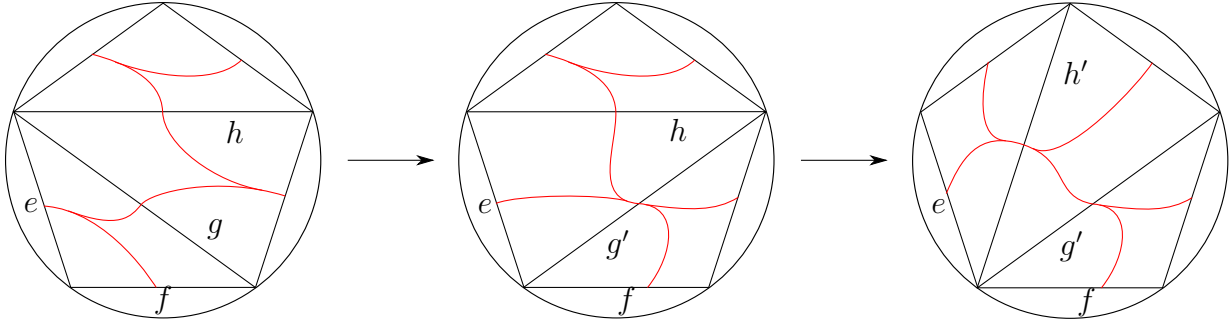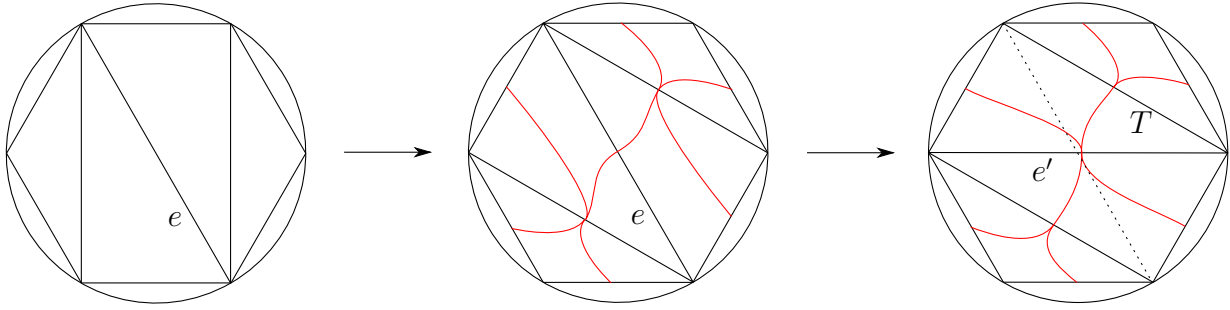

FIGURE 14. Left Whitehead moves must be followed by left moves

Conversely, suppose we have a layered taut ideal triangulation such that the link of every edge is veering. We want to show how to associate to each triangle of the triangulation a train track switch. These switches need to have the property that the induced branched surface on the boundary of a taut tetrahedron T corresponds to a fold, with the front edge e' of T corresponding to the large branch, and the back edge e of T corresponding to the small branch. We note that the train track on the front two faces of T is uniquely determined by the fact that the front edge e' is dual to a large branch. So we choose the switch in each triangle by the choice consistent with the taut tetrahedron in back of the triangle. Now we check that the veering condition implies that this choice is globally consistent.

For each taut tetrahedron T, the switches in the front two faces are determined by our convention so that e' is a large branch, so we need only check what happens on the back two faces of T. Take the two taut tetrahedra adjacent to these back faces. The veering property implies that the back edge e of T is a small branch of the induced train track (see Figure 15). Thus, we see that for any closed loop of Whitehead moves corresponding to the layered taut veering triangulation, we get a closed folding sequence. To finish, we need to see that the folding sequence gives rise to a maximal splitting sequence. The fact that

FIGURE 15. Veering determines the train track on the boundary of T

the folding sequence has an invariant class of projective measures follows from a Perron-Frobenius argument since the induced map of weight spaces is a Perron-Frobenius matrix, as in [19]. Thus, the invariant measure gives rise to a lamination invariant under the monodromy φ. Reversing the folding sequence gives a splitting sequence invariant under φ. Now we apply a result of Hamenstadt which implies that any two splitting sequences are related by commutations [6, Section 5]. The proof of Hamenstadt's result may be made similarly to the proof of Corollary 3.4. This implies that the layered veering triangulation is the same as the one obtained by a maximal splitting sequence. $\square$

Corollary 4.3. *A periodic splitting sequence of a pseudo-Anosov mapping torus coming from $\mathcal{L}^s$ gives rise to a sequence of Whitehead moves, which when reversed corresponds to a periodic splitting sequence associated to $\mathcal{L}^u$.*

Proof. The induced triangulation of the punctured mapping torus is veering. Changing the orientation of all the triangles also gives a veering triangulation. So the reversed sequence of Whitehead moves must also be associated to a periodic splitting sequence, by Proposition 4.2. $\square$

This shows that the layered taut ideal triangulation associated to a pseudo-Anosov map is intrinsic, in that it does not depend on $\mathcal{L}^s$ or $\mathcal{L}^u$.

5. EXAMPLE

For mapping classes of the punctured torus or the 4-punctured sphere, the main construction will produce the canonical layered ideal triangulations considered in [9, 3]. For concreteness, we present the results of the main construction for the case of a 4-strand pseudo-Anosov braid of minimal dilatation. This braid was proven to be minimal dilatation among 4-strand pseudo-Anosov braids by Ko, Los, and Song [21]. The associated pseudo-Anosov map $\varphi : \Sigma_{0,5} \to \Sigma_{0,5}$ has dilatation $\lambda = \lambda(\varphi) = 2.29663\ldots$, where λ is the maximal root of the polynomial $x^4 - 2x^3 - 2x + 1$. The invariant laminations have 5 complementary regions which are punctured monogons, and one complementary region which is a triangle. So $\Sigma_\varphi^\circ = \Sigma_{0,6}$, and the mapping torus $\mathcal{T}(\varphi^\circ)$ is shown in Figure 16. The manifold is a 3-component link complement, which has hyperbolic volume 5.33... (conjecturally this is the minimal volume 3-cusped hyperbolic 3-manifold). Given the lamination data computed by Ko-Los-Song, we found a periodic maximal splitting sequence by hand,
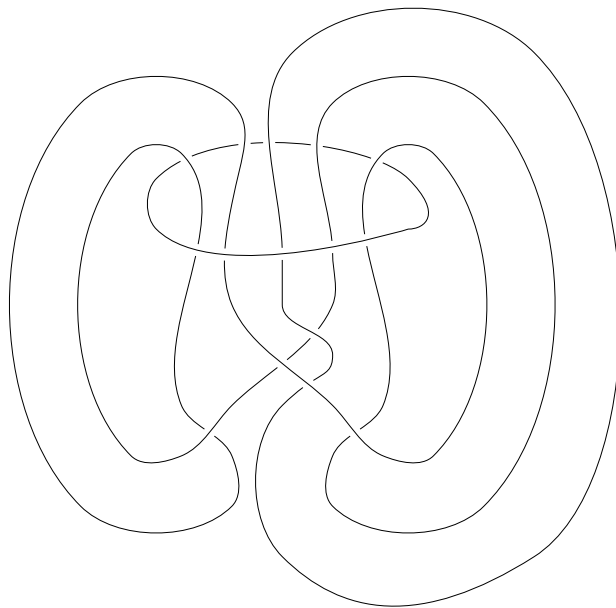
FIGURE 16. Braid representing $\mathcal{T}(\varphi^\circ)$ (the middle strand is the singular fiber)

shown in Figure 17. Since there are six splits, the manifold has a 6 tetrahedron layered
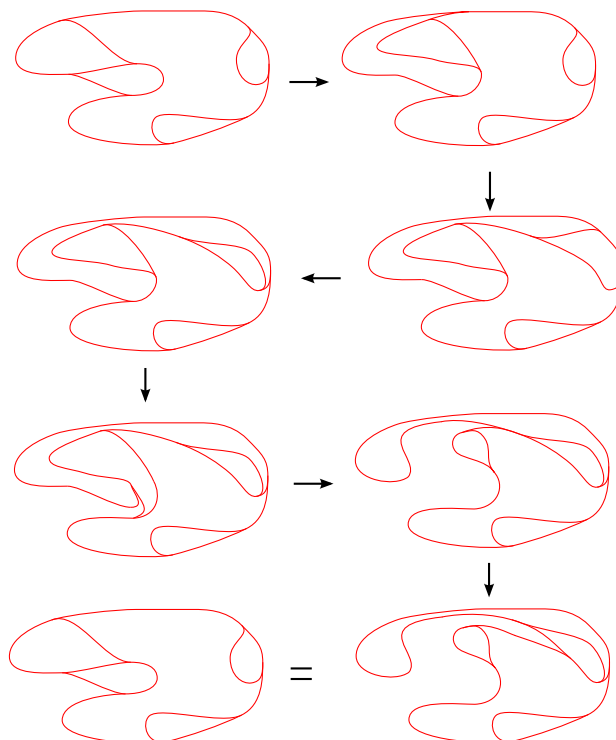

FIGURE 17. A maximal splitting sequence for φ

veering triangulation. If we represent the tetrahedra by ideal hyperbolic simplices, then

they form an ideal hyperbolic tetrahedron triangulation of the manifold, since if one of the tetrahedra had reversed orientation, the volume would be < 5.3.

6. DILATATION BOUNDS

This section gives an alternative approach to the paper [2]. Given a pseudo-Anosov homeomorphism φ of a surface $\Sigma_{g,n}$, $\chi(\Sigma_{g,n}) < 0$, let $\lambda(\varphi)$ denote its dilatation. For any $P > 1$, we define

$$\Psi_P = \left\{ \varphi : \Sigma_{g,n} \to \Sigma_{g,n} \mid \varphi \text{ pseudo-Anosov, and } \lambda(\varphi) \leq P^{\frac{1}{2g-2+2/3n}} \right\}.$$

It follows from work of Penner [19] that for P sufficiently large, and for each closed surface Σ_g of genus $g \geq 2$, there exists $\varphi_g : \Sigma_g \to \Sigma_g$ so that

$$\{\varphi_g : \Sigma_g \to \Sigma_g\}_{g=2}^{\infty} \subset \Psi_P.$$

Given a pseudo-Anosov homeomorphism $\varphi : \Sigma \to \Sigma$, let $\Sigma_\varphi^\circ = \Sigma^\circ \subset \Sigma$ be the surface obtained by removing the singularities of the stable and unstable foliations for φ and let $\varphi|_{\Sigma^\circ} : \Sigma^\circ \to \Sigma^\circ$ denote the restriction. The set of pseudo-Anosov homeomorphisms

$$\Psi_P^\circ = \{\varphi|_{\Sigma^\circ} : \Sigma^\circ \to \Sigma^\circ \mid (\varphi : \Sigma \to \Sigma) \in \Psi_P\}$$

is therefore also infinite.

Lemma 6.1. *A maximal train track on $\Sigma_{g,n}$ has at most $18g - 18 + 6n$ branches.*

Proof. Assume every complementary region of τ is a punctured monogon or a trigon. Then the number of monogons is at most n, and let t denote the number of trigons. Let e denote the number of branches, and let v denote the number of switches of the train track. Since each switch is in one to one correspondence with a cusp of a complementary region, we get $v = n + 3t$. Moreover, $3v = 2e$ counts the number of pairs of incidences between branches and switches. The euler characteristic is given by $\chi(S) = 2 - 2g - n = v - e + t$. Solving for e, using the equations $t = (v - n)/3$ and $v = \frac{2}{3}e$ to eliminate v, t, we get $e = 18g - 18 + 6n$. $\square$

The following theorem gives a refinement of [2, Theorem 1.1].

Theorem 6.2. *Let $M \in \mathcal{T}(\Psi_P)$ be a mapping torus of a punctured pseudo-Anosov class $M = \mathcal{T}(\varphi^\circ)$, $\varphi \in \Psi_P$. Then M has a taut ideal triangulation with at most $\frac{1}{2}(P^9 - 1)$ tetrahedra.*

Proof. Let e be the number of branches of a train track τ fully carrying a stable lamination for φ. By the previous lemma, $e \leq 18g - 18 + 6n$. Then Σ° has an ideal triangulation with e edges. Let $\tau = \sigma_0 \to \sigma_1 \to \cdots \to \sigma_m = \varphi(\tau)$ be an invariant sequence of train track foldings coming from Theorem 3.5. The weight space $W(\sigma_i)$ has positive coordinates in a space of dimension e. Folding changes the weight on the folded edge by the sum of the weights of the two edges folded onto it (see Figure 4). Since there is a 1-1 correspondence between the edges of σ_i and σ_{i+1}, we can think of the map $W(\sigma_i) \to W(\sigma_{i+1})$ as a unipotent matrix, with the sum of entries $e + 2$. Mutiplying such matrices at least adds the off-diagonal entries. So the induced map $W(\tau) \to W(\varphi(\tau))$ is a matrix with the sum of the entries at least $2m + e$. Then from [5, Lemma 3.1], we get $2m + 1 \leq \lambda(\varphi)^e$. By the hypothesis that

$\varphi \in \Psi_P$, we have $\lambda(\varphi) \leq P^{\frac{9}{e}}$, so we conclude that $2m + 1 \leq P^9$. Thus, $\mathcal{T}(\varphi^\circ)$ has a taut ideal layered triangulation with at most $\frac{1}{2}(P^9 - 1)$ tetrahedra. $\square$

In particular, the collection of mapping tori $\{\mathcal{T}(\varphi^\circ) | \varphi \in \Psi_P\}$ is finite. Let's apply this result to the case of the closed surface of genus g. Let δ_g be the minimal dilatation of a pseudo-Anosov map of a closed surface Σ_g of genus g. By a result of Hironaka and Kin, $\delta_g^{g-1} \leq 2 + \sqrt{3}$. Thus, if φ is a pA map of Σ_g with dilatation δ_g, then $\lambda(\varphi) = \delta_g \leq (2+\sqrt{3})^{\frac{1}{g-1}} = (2+\sqrt{3})^{\frac{2}{-\chi(S)}}$. So we may take $P = (2+\sqrt{3})^2$ in the above theorem, and we obtain the number of tetrahedra of a taut ideal triangulation is bounded by $\frac{1}{2}((2+\sqrt{3})^{18} - 1) \leq 10^{10}$.

7. Conclusion

There are several natural questions arising from this paper.

First, in the previous section we gave an upper bound on the number of tetrahedra needed for a taut ideal layered triangulation of the punctured mapping torus of a minimal dilatation pA map of a closed surface Σ_g. One ought to be able to improve this estimate from an improvement of the upper bound on δ_g in the papers [8, 1]. Moreover, one ought to be able to improve the lower bound given in Theorem 6.2 by getting a better estimate of the Perron-Frobenius eigenvalues of the transformations of weight spaces. It is an interesting question to compute the numbers δ_g or more generally $\delta_{g,n}$. The splitting sequences in this paper give a possible approach to this for a given g, n. One may form finitely many automata of folding sequences for a given surface $\Sigma_{g,n}$, one for each possible collection of indices of the singular points. Then one can compute the minimal dilatations for closed paths in these folding automata. Unfortunately, there appears to be many more train tracks which appear in these automata than in those considered by Ko-Los-Song [21, 5].

Secondly, the main construction of a taut ideal layered triangulation associated to a pseudo-Anosov mapping class gives a new classification of conjugacy classes of pA maps. To a pA mapping class, one may associate the taut layered triangulation of the mapping torus, together with an encoding of the homology class of the fiber. This homology class may be represented canonically as a harmonic simplicial 2-cycle for the taut triangulation. This amounts to assigining weights to the faces of the triangulation such that the boundary is zero, and the signed sum of the weights on the faces of a taut tetrahedron is zero (corresponding to the dual 1-chain being a 1-cycle). Then two pA mapping classes are conjugate if and only if there is a combinatorial equivalence between the two triangulations which preserves the 2-cycles. This gives a new way of encoding pseudo-Anosov conjugacy classes, which seems to be somewhat simpler than previous methods (see [7, 12, 13, 14]). There is also a conjugacy invariant implicit in the geometrization theorem for mapping tori [24]. In particular, our approach simplifies Hamenstadt's conjugacy invariant coming from her study of cube complexes of train track splitting sequences [6], but unpublished. It would be interesting to implement this classification algorithmically, and analyze the computational complexity.

Thirdly, it would be interesting to understand the implications of a veering taut ideal triangulation which is not layered. There are two branched surfaces naturally associated to such a triangulation as constructed in the proof of Proposition 4.2, which should carry

essential laminations. It would be interesting to find an example of a veering triangulation which does not come from a fibration (that is, is not layered).

Finally, it is an interesting question whether the veering triangulations studied in Section 4 can be realized in the hyperbolic metric as ideal hyperbolic triangulations with all positively oriented tetrahedra. This is true in the punctured torus case [9, 11, 4], and in the example investigated in Section 5. It's possible that the approach of [4] may be generalized to deal with this question for veering triangulations.

References

[1] John William Aaber and Nathan M. Dunfield, *Closed surface bundles of least volume*, preprint, February 2010, arXiv:1002.3423.

[2] Benson Farb, Chris Leininger, and Dan Margalit, *Small dilatation pseudo-anosovs and 3-manifolds*, 40 pages, May 2009, arXiv:0905.0219.

[3] W. Floyd and A. Hatcher, *Incompressible surfaces in punctured-torus bundles*, Topology Appl. **13** (1982), no. 3, 263–282.

[4] François Guéritaud, *On canonical triangulations of once-punctured torus bundles and two-bridge link complements*, Geom. Topol. **10** (2006), 1239–1284, With an appendix by David Futer.

[5] Ji-Young Ham and Won Taek Song, *The minimum dilatation of pseudo-Anosov 5-braids*, Experiment. Math. **16** (2007), no. 2, 167–179.

[6] Ursula Hamenstädt, *Geometry of the mapping class groups. I. Boundary amenability*, Invent. Math. **175** (2009), no. 3, 545–609.

[7] Geoffrey Hemion, *On the classification of homeomorphisms of 2-manifolds and the classification of 3-manifolds*, Acta Math. **142** (1979), no. 1-2, 123–155.

[8] Eriko Hironaka, *Small dilatation pseudo-anosov mapping classes coming from the simplest hyperbolic braid*, preprint, September 2009, arXiv:0909.4517.

[9] Troels Jørgensen, *On pairs of once-punctured tori*, Kleinian groups and hyperbolic 3-manifolds (Warwick, 2001), London Math. Soc. Lecture Note Ser., vol. 299, Cambridge Univ. Press, Cambridge, 2003, pp. 183–207.

[10] Marc Lackenby, *Taut ideal triangulations of 3-manifolds*, Geom. Topol. **4** (2000), 369–395 (electronic).

[11] ———, *The canonical decomposition of once-punctured torus bundles*, Comment. Math. Helv. **78** (2003), no. 2, 363–384.

[12] Lee Mosher, *Pseudo-anosovs on punctured surfaces*, Ph.D. thesis, Princeton, 1983.

[13] ———, *The classification of pseudo-Anosovs*, Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984), London Math. Soc. Lecture Note Ser., vol. 112, Cambridge Univ. Press, Cambridge, 1986, pp. 13–75.

[14] ———, *Train track expansions of measured foliations*, preprint, 2003, http://andromeda.rutgers.edu/~mosher/.

[15] Athanase Papadopoulos and Robert C. Penner, *A characterization of pseudo-Anosov foliations*, Pacific J. Math. **130** (1987), no. 2, 359–377.

[16] ———, *A characterization of pseudo-Anosov foliations*, Pacific J. Math. **130** (1987), no. 2, 359–377.

[17] ———, *Enumerating pseudo-Anosov foliations*, Pacific J. Math. **142** (1990), no. 1, 159–173.

[18] R. C. Penner and J. L. Harer, *Combinatorics of train tracks*, Annals of Mathematics Studies, vol. 125, Princeton University Press, Princeton, NJ, 1992.

[19] Robert C. Penner, *Bounds on least dilatations*, Proceedings of the American Mathematical Society **113** (1991), no. 2, 443–450.

[20] Valentin Poénaru, *Travaux de Thurston sur les difféomorphismes des surfaces et l'espace de Teichmüller*, Séminaire Bourbaki (1978/79), Lecture Notes in Math., vol. 770, Springer, Berlin, 1980, pp. Exp. No. 529, pp. 66–79.

[21] Won Taek Song, Ki Hyoung Ko, and Jérôme E. Los, *Entropies of braids*, J. Knot Theory Ramifications **11** (2002), no. 4, 647–666, Knots 2000 Korea, Vol. 2 (Yongpyong).

[22] William P. Thurston, *The geometry and topology of 3-manifolds*, Lecture notes from Princeton University, 1978–80, http://www.msri.org/communications/books/gt3m.

[23] William P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) **19** (1988), no. 2, 417–431.

[24] William P. Thurston, *Hyperbolic structures on 3-manifolds, ii: Surface groups and 3-manifolds which fiber over the circle*, preprint, 1998, arXiv:math/9801045.

University of California, Berkeley, 970 Evans Hall #3840, Berkeley, CA 94720-3840
E-mail address: ianagol@math.berkeley.edu